Connect. Accelerate. Outperform.™

MLNX_EN for Linux Release Notes

Rev 3.1-1.0.4

www.mellanox.com

NOTE:
THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT (“PRODUCT(S)”) AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES “AS-IS” WITH ALL FAULTS OF ANY KIND AND SOLELY FOR THE PURPOSE OF AIDING THE CUSTOMER IN TESTING APPLICATIONS THAT USE THE PRODUCTS IN DESIGNATED SOLUTIONS. THE CUSTOMER'S MANUFACTURING TEST ENVIRONMENT HAS NOT MET THE STANDARDS SET BY MELLANOX TECHNOLOGIES TO FULLY QUALIFY THE PRODUCT(S) AND/OR THE SYSTEM USING IT. THEREFORE, MELLANOX TECHNOLOGIES CANNOT AND DOES NOT GUARANTEE OR WARRANT THAT THE PRODUCTS WILL OPERATE WITH THE HIGHEST QUALITY. ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT ARE DISCLAIMED. IN NO EVENT SHALL MELLANOX BE LIABLE TO CUSTOMER OR ANY THIRD PARTIES FOR ANY DIRECT, INDIRECT, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING, BUT NOT LIMITED TO, PAYMENT FOR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY FROM THE USE OF THE PRODUCT(S) AND RELATED DOCUMENTATION EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Mellanox Technologies
350 Oakmead Parkway Suite 100
Sunnyvale, CA 94085
U.S.A.
www.mellanox.com
Tel: (408) 970-3400
Fax: (408) 970-3403

© Copyright 2015. Mellanox Technologies. All Rights Reserved.

Mellanox®, Mellanox logo, BridgeX®, CloudX logo, Connect-IB®, ConnectX®, CoolBox®, CORE-Direct®, GPUDirect®, InfiniHost®, InfiniScale®, Kotura®, Kotura logo, Mellanox Federal Systems®, Mellanox Open Ethernet®, Mellanox ScalableHPC®, Mellanox Connect Accelerate Outperform logo, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, Open Ethernet logo, PhyX®, SwitchX®, TestX®, The Generation of Open Ethernet logo, UFM®, Virtual Protocol Interconnect®, Voltaire® and Voltaire logo are registered trademarks of Mellanox Technologies, Ltd.

Accelio™, CyPU™, FPGADirect™, HPC-X™, InfiniBridge™, LinkX™, Mellanox Care™, Mellanox CloudX™, Mellanox Multi-Host™, Mellanox NEO™, Mellanox PeerDirect™, Mellanox Socket Direct™, Mellanox Spectrum™, NVMeDirect™, StPU™, Spectrum logo, Switch-IB™, Unbreakable-Link™ are trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

Table of Contents

Table of Contents 1
List of Tables 1
Release Update History 2
Chapter 1 Introduction 3
1.1 Supported Platforms and Operating Systems 3
1.2 Supported Hypervisors 4
1.3 Supported NICs Firmware Versions 4
Chapter 2 Changes and New Features in Rev 3.1-1.0.4 5
2.1 Unsupported Functionality/Features 5
Chapter 3 Known Issues 6
3.1 Driver Installation/Loading/Unloading/Start Known Issues 6
3.2 Performance Known Issues 7
3.3 NICs Known Issues 8
3.4 Ethernet Network 9
3.5 Resiliency Known Issues 12
3.6 Virtualization 13
Chapter 4 Bug Fixes History. 14
Chapter 5 Change Log History 18

List of Tables

Table 1: Release Update History 2
Table 2: Supported Uplinks to Servers 3
Table 3: Supported Platforms and Operating Systems 3
Table 4: Supported NICs Firmware Versions 4
Table 5: Changes in v3.1-1.0.4 5
Table 6: Driver Installation Issues 6
Table 7: Driver Start Issues 6
Table 8: Driver Loading Issues 7
Table 9: Driver Start Known Issues 7
Table 10: Performance Known Issues 7
Table 11: ConnectX®-4 (mlx5 Driver) Known Issues 8
Table 12: Ethernet Known Issues 9
Table 13: Quality of Service Issues 11
Table 14: Flow Steering Issues 11
Table 15: Ethernet Performance Counters Known Issues 11
Table 16: Resiliency Known Issues 12
Table 17: SR-IOV Issues 13
Table 18: Fixed Bugs List 14
Table 19: Change Log History 18

Release Update History

Table 1 - Release Update History

Release	Date	Description
Rev 3.1-1.0.4	October 08th, 2015	Initial version

1 Introduction

These are the release notes for Mellanox Technologies' MLNX_EN for Linux Rev 3.1-1.0.4driver kit for Mellanox adapter cards supporting the following uplinks to servers:

Table 2 - Supported Uplinks to Servers

Uplink/NICs	Uplink Speed
ConnectX®-4	• Ethernet: 10GigE, 25GigE, 40GigE, 50GigE and 100GigE
ConnectX®-4 Lx	• Ethernet: 10GigE, 25GigE, 40GigE, and 50GigE
ConnectX®-3/ConnectX®-3 Pro	• Ethernet: 10GigE, 40GigE and 56GigE[a]
ConnectX®-2	• Ethernet: 10GigE, 20GigE
PCI Express 2.0	2.5 or 5.0 GT/s
PCI Express 3.0	8 GT/s

a. 56 GbE is a Mellanox propriety link speed and can be achieved while connecting a Mellanox adapter cards to Mellanox SX10XX switch series or connecting a Mellanox adapter card to another Mellanox adapter card.

For additional information on the new features, please refer to the MLNX_EN User Manual.

1.1 Supported Platforms and Operating Systems

The following are the supported OSs in MLNX_EN for Linux Rev 3.1-1.0.4:

Table 3 - Supported Platforms and Operating Systems

Operating System	Platform
RHEL/CentOS 6.5	x86_64
RHEL/CentOS 6.6	x86_64
RHEL/CentOS 6.7	x86_64/PPC64 (Power 7)
RHEL/CentOS 7.0	x86_64PPC64 (Power 7)
RHEL/CentOS 7.1	x86_64/PPC64 (Power 7)/PPC64le (Power 8)/ ARM64 (ARM is at beta level)
SLES11 SP1	x86_64
SLES11 SP2	x86_64
SLES11 SP3	x86_64/PPC64 (Power 7)
SLES11 SP4	x86_64/PPC64 (Power 7)
SLES12	x86_64/PPC64le (Power 8)
OEL 6.3	x86_64
OEL 6.4	x86_64
OEL 6.5	x86_64
OEL 6.6	x86_64
OEL 6.7	x86_64
OEL 7.0	x86_64
OEL 7.1	x86_64
Fedora 19	x86_64/PPC64 (Power 7)

Table 3 - Supported Platforms and Operating Systems

Operating System	Platform
Fedora 20	x86_64
Fedora 21	x86_64/PPC64 (Power 7)
Ubuntu 12.04	x86_64
Ubuntu 14.04	x86_64/PPC64le (Power 8)
Ubuntu 14.10	x86_64/PPC64le (Power 8)
Ubuntu 15.04	x86_64/PPC64le (Power 8)
Debian 6.0.10	x86_64
Debian 7.6	x86_64
Debian 8.0	x86_64
Debian 8.1	x86_64
kernel 3.10[a] - 4.1	

a. This kernel is supported only when using the Operating Systems stated in the table above.

1.2 Supported Hypervisors

The following are the supported hypervisors in MLNX_EN Rev 3.1-1.0.4:

- KVM:
  - RedHat 6.6, 6.7, 7.1
  - Ubuntu 14.10, 15.04
  - Sles11SP4, Sles12
  - Debian 6.0.10
- Xen4.2
- XenServer6.5

1.3 Supported NICs Firmware Versions

MLNX_EN Rev 3.1-1.0.4 supports the following Mellanox network adapter cards firmware versions:

Table 4 - Supported NICs Firmware Versions

NIC	Recommended Firmware Rev.	Additional Firmware Rev. Supported
ConnectX®-4 Lx	14.12.1100	-
ConnectX®-4	12.12.1100	12.12.0780
ConnectX®-3 Pro	2.35.5100	2.34.5000
ConnectX®-3	2.35.5100	2.34.5000
ConnectX®-2	2.9.1000	2.9.1000

For official firmware versions please see:
http://www.mellanox.com/content/pages.php?pg=firmware_download

2 Changes and New Features in Rev 3.1-1.0.4

Table 5 - Changes in v3.1-1.0.4

Category	Description
Wake-on-LAN (WOL)	Wake-on-LAN (WOL) is a technology that allows a network professional to remotely power on a computer or to wake it up from sleep mode.
Hardware Accelerated 802.1ad VLAN (Q-in-Q Tunneling)	Q-in-Q tunneling allows the user to create a Layer 2 Ethernet connection between two servers. The user can segregate a different VLAN traffic on a link or bundle different VLANs into a single VLAN.
ConnectX-4 ECN	ECN in ConnectX-4 enables end-to-end congestions notifications between two end-points when a congestion occurs, and works over Layer 3.
Minimal Bandwidth Guarantee (ETS)	The amount of bandwidth (BW) left on the wire may be split among other TCs according to a minimal guarantee policy.
SR-IOV Ethernet	SR-IOV Ethernet in ConnectX-4 at Beta level

2.1 Unsupported Functionality/Features

The following are the unsupported functionalities/features in MLNX_EN Rev 3.1-1.0.4:

- pm_qos API - Power Management
- Adaptive Interrupt Moderation Algorithm
- Virtual Guest Tagging (VGT+)

3 Known Issues

The following is a list of general limitations and known issues of the various components of this MLNX_EN for Linux release.

3.1 Driver Installation/Loading/Unloading/Start Known Issues

3.1.1 Driver Installation Known Issues

Table 6 - Driver Installation Issues

Index	Description	Workaround
1.	Firmware update fails during mlx4_en standalone installation on guest machine connect to NIC in Passthrough mode.	-
2.	mlx4_en standalone drivers are not available after installation on SLSE11 SP3 PPC64.	Run the installation with the "--disable-kmp" flag.
3.	On RHEL 7, MLNX_EN driver v2.3-1.0.0 installation fails the first time. However, the driver is installed successfully on the second attempt.	To avoid failure on the first attempt, prior to running the installation, run: `make -C /usr/src/kernels/$(uname -r) kernelrelease`

3.1.2 Driver Start Known Issues

Table 7 - Driver Start Issues

Index	Description	Workaround
1.	Driver restart fails after mlnx_en installation if MLNX_OFED was previously installed on the system.	1. Reboot the server after mlnx_en installation. 2. Answer **'n'** to the following question at the end of the installation: `"Do you wish to reload the driver now? (y/n) [y] "` n 3. Run `"apt-get remove --purge mlnx-ofed-kernel-utils"` 4. Reboot the server after mlnx_en installation.
2.	In ConnectX®-2, if the driver load succeeds, the informative message below is presented conveying the below limitations: `If port type is ETH then the maximum priority for VLAN tagged is 3 mlx4_core 0000:0d:00.0: command SET_PORT (0xc) failed: in_param=0x120064000, in_mod=0x2, op_mod=0x0, fw status = 0x40`	-

3.1.3 Driver Loading Known Issues

Table 8 - Driver Loading Issues

Index	Description	Workaround
1.	On Debian 7.x, the Inbox driver is loaded instead of mlx4_en standalone drivers	Add the following lines to `"/etc/rc.local"`: • `/sbin/modprobe -r mlx-4_fc > /dev/null 2>&1` • `/sbin/modprobe -r mlx-4_en > /dev/null 2>&1` • `/sbin/modprobe -r mlx-4_ib > /dev/null 2>&1` • `/sbin/modprobe -r mlx-4_core > /dev/null 2>&1` • `# load the new driver` • `/sbin/modprobe mlx4_en > /dev/null 2>&1`

3.1.4 Driver Start Known Issues

Table 9 - Driver Start Known Issues

Index	Description	Workaround
1.	"Out of memory" issues may rise during drivers load depending on the values of the driver module parameters set (e.g. log_num_cq)	-

3.2 Performance Known Issues

Table 10 - Performance Known Issues

Index	Description	Workaround
1.	On machines with irqbalancer daemon turned off, the default Ethernet interrupts will be routed to a single core which may cause overload and software/ hardware lockups.	Execute the following script as root: `set_irq_affinity.sh <interface or Eth device> [2nd interface or Eth device]`
2.	Out of the box throughput performance in Ubuntu14.04 is not optimal and may achieve results below the line rate in 40GE link speed.	For additional performance tuning, please refer to Performance Tuning Guide.
3.	UDP receiver throughput may be lower then expected, when running over mlx4_en Ethernet driver. This is caused by the adaptive interrupt moderation routine, which sets high values of interrupt coalescing, causing the driver to process large number of packets in the same interrupt, leading UDP to drop packets due to overflow in its buffers.	Disable adaptive interrupt moderation and set lower values for the interrupt coalescing manually. `ethtool -C <eth>X adaptive-rx off rx-usecs 64 rx-frames 24` Values above may need tuning, depending the system, configuration and link speed.
4.	Performance degradation might occur when bonding Ethernet interfaces on Centos 6.5	-

3.3 NICs Known Issues

3.3.1 ConnectX®-4 (mlx5 Driver) Known Issues

Table 11 - ConnectX®-4 (mlx5 Driver) Known Issues

Index	Description	Workaround
1.	EEH events that arrive while the mlx5 driver is loading may cause the driver to hang.	-
2.	The mlx5 driver can handle up to 5 EEH events per hour.	If more events are received, cold reboot the machine.
3.	In PPC systems, when working with ConnextX®-4 adapter card configured as Ethernet, driver load fails with BAD INPUT LENGTH. dmesg: `command failed, status bad input length(0x50), syndrome 0x9074aa`	-
4.	Changing the link speed is not supported in Ethernet driver when connected to a ConnectX-4 card.	-
5.	Changing the RX queues number is not supported in Ethernet driver when connected to a ConnectX-4 card.	-
6.	Error counters such as: CRC error counters, RX out range length error counter, are missing in the ConnectX-4 Ethernet driver.	-

3.4 Ethernet Network

3.4.1 Ethernet Known Issues

Ethernet Know Issues are applicable to ConnectX-3/ConnectX-3 Pro only.

Table 12 - Ethernet Known Issues

Index	Description	Workaround
1.	Ethernet PV VLAN Guest transparent Tagging (VGT) is only supported in openvswitch and not in standard Linux vBridges and libvirt For more information please see : http://libvirt.org/formatnetwork.html (Setting VLAN tag section)	-

2.	Changing the ring size on 32-bit system may result in failure due to lack of memory. Therefore, mlx4_en will not be able to vmap enough memory and the below message will be printed in dmesg:
	`vmap allocation for size 528384 failed: use vmalloc=<size> to increase size`
	`In this case user can enlarge the vmalloc memory by adding vmalloc=<size> to grub.conf`
	`Default vmalloc setting is 128M. It is recommended to add each time 64M of memory until desired ring size can be allocated.`
	Please note, that in case vmalloc size is too big, the OS will fail to boot, so please use caution when adding additional memory.
	For more info refer to: http://www.mythtv.org/wiki/Common_Problem:_vmalloc_too_small
		-

3.	Kernel panic might occur during fio splice in kernels before 2.6.34-rc4.	Use kernel v2.6.34-rc4 which provides the following solution: `baff42a net: Fix oops from tcp_collapse() when using splice()`
4.	When creating more than 125 VLANs and SR-IOV mode is enabled, a kernel warning message will be printed indicating that the native VLAN is created but will not work with RoCE traffic. kernel warning: `mlx4_core 0000:07:00.0: vhcr command ALLOC_RES (0xf00) slave:0 in_param 0x7e in_mod=0x107, op_mod=0x1 failed with error:0, status -28`	
5.	In PPC systems when QoS is enabled a harmless Kernel DMA mapping error messages might appear in kernel log (iommu related issue).	-

Table 12 - Ethernet Known Issues (Continued)

Index	Description	Workaround
6.	The default priority to TC mapping assigns all priorities to TC0. This configuration achieves fairness in transmission between priorities but may cause undesirable PFC behavior where pause request for priority "n" affects all other priorities.	Run: `mlnx_qos -i <dev> -p 0,1,2,3,4,5,6,7 -s ets,ets,ets,ets,ets,ets,ets,ets -t 12,13,12,13,12,13,12,13` This needs to be applied every time after loading the `mlx4_en` driver.
7.	Transmit timeout might occur on RH6.3 as a result of lost interrupt (OS issue). In this case, the following message will be shown in dmesg: `do_IRQ: 0.203 No irq handler for vector (irq -1)`	-
8.	In ConnectX®-2, RoCE UD QP does not include VLAN tags in the Ethernet header	
9.	VXLAN may not be functional when configured over Linux bridge in RH7.0 or Ubuntu14.04. The issue is within the bridge modules in those kernels. In Vanilla kernels above 3.16 issues is fixed.	-
10.	In RH6.4, ping may not work over VLANs that are configured over Linux bridge when the bridge has a mlx4_en interface attached to it.	-
11.	The interfaces LRO needs to be set to "OFF" manually when there is a bond configured on Mellanox interfaces with a Bridge over that bond.	Run: `ethtool -K ethX lro off`
12.	On SLES12, the bonding interface over Mellanox Ethernet slave interfaces does not get IP address after reboot.	1. Set "STARTMODE=hotplug" in the bonding slave's ifcfg files. More details can be found in the SUSE documentations page: https://www.suse.com/documentation/sles-12/book_sle_admin/?page=/documentation/sles-12/book_sle_admin/data/sec_bond.html 2. Enable the "nanny" service to support hot-plugging: Open the `"/etc/wicked/common.xml"` file. Change: `"<use-nanny>false</use-nanny>"` to `"<use-nanny>true</use-nanny>"` 3. Run: `# systemctl restart wickedd.service wicked`
13.	ethtool -x command does not function in SLES OS.	-
14.	Ethertype proto 0x806 not supported by ethtool.	-
15.	ETS is not supported in kernels that do not have MQPRIO as QDISC_KIND option in the tc tool.	-

Table 12 - Ethernet Known Issues (Continued)

Index	Description	Workaround
16.	When NC-SI is ON, the port's MTU cannot be set to lower than 1500.	-
17.	GRO is not functional when using VXLAN in ConnectX-3 adapter cards.	-
18.	ethtool -X: The driver supports only the `'equal'` mode and cannot be set by using weight flags.	-
19.	Q-in-Q infrastructure in the kernel is supported only from kernel version 3.10 and up.	-
20.	When SLES11 SP4 is used as a DHCP client over ConnectX-3 or ConnectX-3 adapters, it might fail to get an IP from the DHCP server.	-

3.4.2 Quality of Service Known Issues

Table 13 - Quality of Service Issues

Index	Description	Workaround
1.	QoS is not supported in XenServer, Debian 6.0 and 6.2 with uek kernel	-
2.	When QoS features are not supported by the kernel, `mlnx_qos` tool may crash.	-

3.4.3 Flow Steering Known Issues

Table 14 - Flow Steering Issues

Index	Description	Workaround
1.	Flow Steering is disabled by default in firmware version < 2.32.5100.	To enable it, set the parameter below as follow: `log_num_mgm_entry_size should set to -1`
2.	IPv4 rule with source IP cannot be created in SLES 11.x OSes.	-
3.	RFS does not support UDP.	-

3.4.4 Ethernet Performance Counters Known Issues

Table 15 - Ethernet Performance Counters Known Issues

Index	Description	Workaround
1.	In a system with more than 61 VFs, the 62nd VF and onwards is assigned with the SINKQP counter, and as a result will have no statistics, and loopback prevention functionality for SINK counter.	-

Table 15 - Ethernet Performance Counters Known Issues (Continued)

Index	Description	Workaround
2.	Since each VF tries to allocate 2 more QP counter for its RoCE traffic statistics, in a system with less than 61 VFs, if there is free resources it receives new counter otherwise receives the default counter which is shared with Ethernet. In this case RoCE statistics is not available.	-
3.	In ConnectX®-3, when we enable function-based loopback prevention for Ethernet port by default (i.e., based on the QP counter index), the dropped self-loopback packets increase the IfRxErrorFrames/Octets counters.	-

3.5 Resiliency Known Issues

3.5.1 Reset Flow Known Issues

Table 16 - Resiliency Known Issues

Index	Description	Workaround
1.	Reset Flow can run on XenServer 6 only after the active user space applications running verbs are terminated.	-
2.	SR-IOV non persistent configuration (such as VGT, VST, Host assigned GUIDs, and QP0-enabled VFs) may be lost upon Reset Flow.	Reset Admin configuration post Reset Flow
3.	Upon Reset Flow or after running restart driver, Ethernet VLANs are lost.	Reset the VLANs using the `ifup` command.
4.	Restarting the driver or running `connectx_port_config` when Reset Flow is running might result in a kernel panic	-
5.	Networking configuration (e.g. VLANs, IPv6) should be statically defined in order to have them set after Reset Flow as of after restart driver.	-
6.	The qemu version (2.0) provided inbox with Ubuntu 14.04 does not work properly when more than 2 VMs are run over an Ubuntu 14.04 Hypervisor.	-
7.	Attempting to attach a PF to a VM when SR-IOV is already enabled on that PF may result in a kernel panic.	-
8.	On various combinations of Hypervisor/OSes and Guest/OSes, an issue might occur when attaching/detaching VFs to a guest while that guest is up and running.	Attach/detach VFs to/from a VM only while that VM is down.

3.6 Virtualization

3.6.1 SR-IOV Known Issues

Table 17 - SR-IOV Issues

Index	Description	Workaround
1.	Enabling SR-IOV requires appending the `intel_iommu=on` option to the relevant OS in file `/boot/grub/grub.conf/`. Without that SR-IOV cannot be loaded.	-
2.	When using legacy VMs with MLNX_EN 2.x hypervisor, you may need to set the `'enable_64b_cqe_eqe'` parameter to zero on the hypervisor. It should be set in the same way that other module parameters are set for mlx4_core at module load time. For example, add `"options mlx4_core enable_64b_cqe_eqe=0"` as a line in the file `/etc/modprobe.d/mlx4_core.conf`.	-
3.	Attaching or detaching a Virtual Function on SLES11 SP3 to a guest Virtual Machine while the `mlx4_core` driver is loaded in the Virtual Machine may cause a kernel panic in the hypervisor.	Unload the `mlx4_core` module in the hypervisor before attaching or detaching a function to or from the guest.
4.	Enabling SR-IOV requires appending the `"intel_iommu=on"` option to the relevant OS in file `/boot/grub/grub.conf/`. Without that SR-IOV cannot be loaded.	-
5.	The known PCI BDFs for all VFs in kernel command line should be specified by adding `xen-pciback.hide` For further information, please refer to http://wiki.xen.org/wiki/Xen_PCI_Passthrough	-
6.	The qemu version (2.0) provided inbox with Ubuntu 14.04 does not work properly when more than 2 VMs are run over an Ubuntu 14.04 Hypervisor.	-
7.	Attempting to attach a PF to a VM when SR-IOV is already enabled on that PF may result in a kernel panic.	-
8.	On various combinations of Hypervisor/OSes and Guest/OSes, an issue might occur when attaching/detaching VFs to a guest while that guest is up and running.	Attach/detach VFs to/from a VM only while that VM is down.
9.	On ConnectX®-3 NICs with firmware version 2.32.5000 and later, SR-IOV VPI mode works only with Port 1 = ETH and Port 2 = IB.	-
10.	Occasionally, the `lspci | grep Mellanox` command shows incorrect or partial information due to the current `pci.ids` file on the machine.	1. Locate the file: `$locate pci.ids` 2. Manually update the file according to the latest version available online at: https://pci-ids.ucw.cz/v2.2/pci.ids This file can also be downloaded.
11	SR-IOV is not supported in XenServer 6.5.	-
12	SR-IOV is not supported in AMD architecture.	-

4 Bug Fixes History.

Table 18 - Fixed Bugs List

#	Description	Discovered in Release	Fixed in Release
1.	Fixed VXLAN functionality issues.	3.0-2.0.1	3.1-1.0.4
2.	TCP/UDP latency on ConnectX®-4 was higher than expected.	3.0-2.0.1	3.1-1.0.4
3.	TCP throughput on ConnectX®-4 achieved full line rate.	3.0-2.0.1	3.1-1.0.4
4.	Fixed an issue causing inconsistent performance with ConnectX-3 and PowerKVM 2.1.1.	3.0-2.0.1	3.1-1.0.4
5.	Updated the desired `num_entries` in each iteration, and accordingly updated the offset of the WC in the given WC array.	3.0-1.0.1	3.1-1.0.4
6.	Fixed an issue which prevented openibd from starting correctly during boot.	3.0-2.0.1	3.1-1.0.4
7.	Added a new module parameter to control the number of IRQs allocated to the device.	3.0-2.0.1	3.1-1.0.4
8.	Fixed an issue on PPC servers which prevented PCI from reloading after EEH error recovery.	3.0-2.0.1	3.1-1.0.4
9.	Added the option to toggle LRO ON/OFF using the `"-K"` flags. The priv flag hw_lro will determine the type of LRO to be used, if the flag is ON, the hardware LRO will be used, otherwise the software LRO will be used.	3.0-2.0.1	3.1-1.0.4
10.	Fixed race when updating counters.	3.0-2.0.1	3.1-1.0.4
11.	Fixed scheduling while sending atomic dmesg warning during bonding configuration.	3.0-2.0.1	3.1-1.0.4
12.	Added `set_rx_csum` callback implementation.	3.0-2.0.1	3.1-1.0.4
13.	Removed BUG_ON assert when checking if the ring is full.	3.0-1.0.1	3.1-1.0.4
14.	LRO fixes and improvements for jumbo MTU.	2.3-2.0.1	2.4-1.0.0.1
15.	Fixed a crash occurred when changing the number of rings (ethtool set-channels) when interface connected to netconsole.	2.2-1.0.1	2.4-1.0.0.1
16.	Fixed ping issues with IP fragmented datagrams in MTUs 1600-1700.	2.2-1.0.1	2.4-1.0.0.1
17.	The default priority to TC mapping assigns all priorities to TC0. This configuration achieves fairness in transmission between priorities but may cause undesirable PFC behavior where pause request for priority `"n"` affects all other priorities.	2.3-1.0.1	2.4-1.0.0.1
18.	Fixed a crashing issue on non-SR-IOV systems related to VxLAN with dynamic port configuration.	2.3-1.0.0	2.3-2.0.1
19.	Fixed "sleeping while atomic" error occurred when the driver ran many firmware commands simultaneously.	2.3-1.0.0	2.3-2.0.1

Table 18 - Fixed Bugs List

#	Description	Discovered in Release	Fixed in Release
20.	Fixed dmesg warnings: `"NOHZ: local_softirq_pending 08"`.	2.3-1.0.0	2.3-2.0.1
21.	Fixed erratic report of hardware clock which caused bad report of PTP hardware Time Stamping.	2.3-1.0.0	2.3-2.0.1
22.	Fixed counter index allocation for VFs which enables Ethernet port statistics.	2.3-1.0.0	2.3-2.0.1
23.	Fixed kernel panic on Debian-6.0.7 which occurred when the number of TX channels was set above the default value	2.1-1.0.0	2.2-1.0.1
24.	Fixed wrong calculation of packet true-size reporting in LRO flow	2.1-1.0.0	2.2-1.0.1
25.	Fixed a crash incidence which occurred when enabling Ethernet Time-stamping and running VLAN traffic	2.0-3.0.0	2.2-1.0.1
26.	Fixed ALB bonding mode failure when enslaving Mellanox interfaces	2.0-3.0.0	2.1-1.0.0
27.	Fixed leak of mapped memory	2.0-3.0.0	2.1-1.0.0
28.	Fixed TX timeout issues	1.5.10	2.0.-3.0.0
29.	Fixed Ethtool status report for Virtual Functions	1.5.10	2.0.-3.0.0
30.	Fixed an issue of VLAN traffic over Virtual Machine in para-virtualized mode	1.5.10	2.0.-3.0.0
31.	Fixed Ethtool operation crash while interface down	1.5.10	2.0.-3.0.0
32.	Fixed driver unloading after installation. Currently, the driver is reloaded after installation	1.5.9	1.5.10
33.	Fixed an issue related to the configuration files. Configuration files are no longer modified.	1.5.9	1.5.10
34.	Fixed MTU size calculation	1.5.9	1.5.10
35.	Fixed race in join/leave multicast group flow	1.5.9	1.5.10
36.	Restored RX buffer pointer in case of failure	1.5.9	1.5.10
37.	Fixed IPMI	1.5.9	1.5.10
38.	Fixed BlueFlame race	1.5.9	1.5.10
39.	Added sanity check for module parameters	1.5.9	1.5.10
40.	Fixed lost of connectivity bug for iperf long run	1.5.9	1.5.10
41.	Fixed bad TX completion notification request logic, could lead to TX timeout	1.5.8.3	1.5.9
42.	Fixed PowerPC connectivity loss	1.5.8.3	1.5.9
43.	Fixed possible transmit timeout under heavy traffic load	1.5.8.2	1.5.8.3
44.	Fixed error flows in case of MSI-X assignment failures	1.5.7.2	1.5.8.2

Table 18 - Fixed Bugs List

#	Description	Discovered in Release	Fixed in Release
45.	Fixed loss of multicast traffic during addition of multicast addresses	1.5.7.2	1.5.8.2
46.	Enabled legal MTU configuration prior to interface bring-up	1.5.7.2	1.5.8.2
47.	Fixed WoL configuration behavior for single ported devices	1.5.7	1.5.7.2
48.	Fixed the assignment of valid flow control parameters during interface initialization	1.5.7	1.5.7.2
49.	Fixed qdisc behavior when transmitting packets with VLAN ID 0	1.5.7	1.5.7.2
50.	Fixed badly placed counters on ethool -S report	1.5.7	1.5.7.2
51.	Fixed the crash occurred during initialization if an invalid mac address was burned on the card	1.5.6	1.5.7
52.	Fixed a bug where interface link state was not updated correctly in certain scenarios	1.5.6	1.5.7
53.	Fixed a crash occurred when LRO enabled via ethtool but disabled in module parameter	1.5.6	1.5.7
54.	Fixed a warning when using netconsole	1.5.6	1.5.7
55.	Fixed a crash occurred during driver initialization due to multicast list update	1.5.6	1.5.7
56.	Fixed a bug that caused the driver not to work with MSI-X on RedHat EL4 systems	1.5.1.3	1.5.6
57.	Fixed the driver functionality in case the device is enslaved to bond with VLANs over it	1.5.1.3	1.5.6
58.	Fixed the crash occurred when sending malformed TCP packets	1.5.1.3	1.5.6
59.	If the two ports of a ConnectX/ConnectX-2 device are bonded together and one port is closed, then the other port will cease carrying traffic. Fixed in this release.	1.4.2	1.5.1.3
60.	Multiple error messages when working under bonding on some operating systems	1.4.2	1.5.1.3
61.	Fixed a bug avoiding the driver to work with MSI-X on RedHat EL4 systems		1.4.2
62.	Enabled the driver to load on systems with 32 cores and higher		1.4.2
63.	The driver is being stuck if the HW/FW stops responding, reset is done instead		1.4.2
64.	Fixed recovery flows from memory allocation failures		1.4.2
65.	Corrected checksum reports on SLES11 systems		1.4.2
66.	When the system is low on memory, the driver now allocates smaller RX rings		1.4.2

Table 18 - Fixed Bugs List

#	Description	Discovered in Release	Fixed in Release
67.	The driver now retries to obtain MSI-X vectors if the initial request is rejected by the OS		1.4.2

5 Change Log History

Table 19 - Change Log History

Release	Category	Description
3.0-1.0.1	NICs	Added support for ConnectX®-4 Single/Dual-Port Adapter supporting up to 100Gb/s.
	Ignore Frame Check Sequence (FCS) Errors	Upon receiving packets, the packets go through a checksum validation process for the FCS field. If the validation fails, the received packets are dropped. Using this feature, enables you to choose whether or not to drop the frames in case the FCS is wrong and use the FCS field for other info.
	ethtool	Updated ethtool to incorporate ConnectX®-4 adapter card functionalities.
Rev 2.4-1.0.0.1	mlx4_en net-device Ethtool	Added support for Ethtool speed control and advertised link mode.
		Added ethtool txvlan control for setting ON/OFF hardware TX VLAN insertion: `ethtool -k txvlan [on/off]`
		Ethtool report on port parameters improvements.
		Ethernet TX packet rate improvements.
	Virtualization	VXLAN fixes and performance improvements.
	Ethernet net-device	New adaptive interrupt moderation scheme to improve CPU utilization.
Rev 2.3-2.0.1	Reset Flow	Added support for Enhanced Error Handling for PCI (EEH), a recovery strategy for I/O errors that occur on the PCI bus.
	VXLAN	Bug Fixes, see Section 4, “Bug Fixes History.”, on page 14
Rev 2.3-1.0.0	Ethernet	Added support for arbitrary UDP port for VXLAN. From upstream 3.15-rc1 and onward, it is possible to use arbitrary UDP port for VXLAN. This feature requires firmware version 2.32.5100 or higher. Additionally, the following kernel configuration option `CONFIG_MLX4_EN_VXLAN=y` must be enabled.
		MLNX_EN no longer changes the OS sysctl TCP parameters.

Table 19 - Change Log History

Release	Category	Description
Rev 2.2-1.0.1	Reset Flow	Reset Flow is not activated by default. It is controlled by the mlx4_core`'internal_err_reset'` module parameter.
	Ethernet	Ethernet VXLAN support for kernels 3.12.10 or higher
		Power Management Quality of Service: when the traffic is active, the Power Management QoS is enabled by disabling the CPU states for maximum performance.
		Ethernet PTP Hardware Clock support on kernels/OSes that support it
	Performance	Out of the box performance improvements: • Use of affinity hints (based on NUMA node of the device) to indicate the IRQ balancer daemon on the optimal IRQ affinity • Improvement in buffers allocation schema (based on the hint above) • Improvement in the adaptive interrupt moderation algorithm
Rev 2.1-1.0.	mlnx_en	Added reporting autonegotiation support
		Added Transmit Packet Steering (XPS) support
		Added reporting 56Gbit/s link speed support
		Added Receive Flow Steering (RFS) support in UDP
		Added Low Latency Socket (LLS) support
		Added check for dma_mapping errors
Rev 2.0-3.0.0	Operating Systems	Additional OS support: • SLES11SP3 • Fedora16, Fedora17
	Hardware	Added ConnectX-3 Pro support
Rev 1.5.10	General	Section 4, “Bug Fixes History.”, on page 14
Rev 1.5.9	Operating Systems	Added support for kernel.org 3.5
	Performance	Improved latency by optimizing RX repost mechanism
Rev 1.5.8.3	Operating Systems	Added support for RHEL6.3
Rev 1.5.8.2	Operating Systems	Added support for new kernels: 3.1, 3.2, 3.3
Rev 1.5.8.2	Performance	Moved to interrupt mode to handle TX completions
		Added IRQ affinity control scripts (please see README file for more details)
		Optimized Numa aware memory allocations
		Optimized interrupt usage for TX/RX completions
	Installation	Added KMP compliant installation process
	Linux Tools	Added support for Ethtool

Table 19 - Change Log History

Release	Category	Description
Rev 1.5.7.2	Operating Systems	Added support for new OS's: • RHEL6.2 • RHEL5.8 • SLES11SP2
	Performance	Added recording RX queue for GRO packets
		Added the usage of Toeplitz hash function for RSS calculation
	Reports/Statistics	Enabled RXHASH report on supported systems
Rev 1.5.7	Operating Systems	Added support for new OS's: • RHEL6.1 • RHEL5.5 • RHEL5.7 • kernel.org (2.6.37, 2.6.38, 2.6.39, 3.0) • RHEL6.1 KVM
	Performance	Improved performance on PPC systems (Using GRO where LRO is not efficient)
		Added IPv6 support to LRO
		Incremented number of TX and RX queues
		Enabled NAPI usage at any given time
		Enabled TX completions spread among multiple MSI-X vectors
		Improved small packets packet rate
		Added 40GigE support (including Ethtool report)
		Added NUMA support
		Added general performance improvements
Rev 1.5.6	Operating Systems	Added support for new OS's: • RHEL6.0 • RHEL5.6 • SLES11SP1 • kernel.org (2.6.35, 2.6.36)

Table 19 - Change Log History

Release	Category	Description
Rev 1.5.6	Performance	Added blue flame support for kernels > 2.6.28 (improves TX latency by 0.4 usec)
		Added RX acceleration feature that supports recvmsg and recvmmsg system calls. See MLNX_EN_Linux_README for further details.
		Added option to use interrupts for TX completion (polling is the default)
		Added option to disable NAPI (enabled by default)
		Added support for control number of RX rings from module parameter
		Added interrupt vector per each RX ring. See /proc/interrupts
		Adaptive moderation improvements
		Added system tuning option to achieve better performance (idle loop polling)
	Linux Tools	Added hardware revision report via Ethtool
	Multicast Filtering	Added exact match multicast filtering
	Driver Load	Link is brought up upon driver load
Rev 1.5.1.3	Operating Systems	Added support for new OS's: • RHEL5.5 • kernel.org (2.6.16 - 2.6.32)
	Performance	Added UDP RSS support (on ConnectX-2 HW only)
		Improved VLAN tagging performance
	Linux Tools	Ethtool -e support